

RESEARCH ARTICLE

A Statistical Model for *In Vivo* Neuronal Dynamics

Simone Carlo Surace^{1,2*}, Jean-Pascal Pfister²

1 Department of Physiology, University of Bern, Bern, Switzerland, **2** Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland

* surace@ini.uzh.ch



OPEN ACCESS

Citation: Surace SC, Pfister J-P (2015) A Statistical Model for *In Vivo* Neuronal Dynamics. PLoS ONE 10(11): e0142435. doi:10.1371/journal.pone.0142435

Editor: Michele Giugliano, University of Antwerp, BELGIUM

Received: April 15, 2015

Accepted: October 21, 2015

Published: November 16, 2015

Copyright: © 2015 Surace, Pfister. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Intracellular datasets 1, 2, 3 are available at <http://datadryad.org/> under DOI: [10.5061/dryad.k3f5p](https://doi.org/10.5061/dryad.k3f5p). Dataset 4 is available from Haider B, Häusser M, Carandini M at <http://www.carandiniilab.net/awakespontaneousvm>.

Funding: This work was supported by Swiss National Science Foundation: S.C.S. and J.P.P grant PZ00P3_137200 and PP00P3_150637. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing Interests: The authors have declared that no competing interests exist.

Abstract

Single neuron models have a long tradition in computational neuroscience. Detailed biophysical models such as the Hodgkin-Huxley model as well as simplified neuron models such as the class of integrate-and-fire models relate the input current to the membrane potential of the neuron. Those types of models have been extensively fitted to *in vitro* data where the input current is controlled. Those models are however of little use when it comes to characterize intracellular *in vivo* recordings since the input to the neuron is not known. Here we propose a novel single neuron model that characterizes the statistical properties of *in vivo* recordings. More specifically, we propose a stochastic process where the subthreshold membrane potential follows a Gaussian process and the spike emission intensity depends nonlinearly on the membrane potential as well as the spiking history. We first show that the model has a rich dynamical repertoire since it can capture arbitrary subthreshold autocovariance functions, firing-rate adaptations as well as arbitrary shapes of the action potential. We then show that this model can be efficiently fitted to data without overfitting. We finally show that this model can be used to characterize and therefore precisely compare various intracellular *in vivo* recordings from different animals and experimental conditions.

Introduction

During the last decade, there has been an increasing number of studies providing intracellular *in vivo* recordings. From the first intracellular recordings performed in awake cats [1, 2] to more recent recording in cats [3], monkeys [4], mice [5], and even in freely behaving rats [6], it has been shown that the membrane potential displays large fluctuations and is very rarely at the resting potential. Some recent findings in the cat visual cortex have also suggested that the statistical properties of spontaneous activity is comparable to the neuronal dynamics when the animal is exposed to natural images [7]. Similar results have been found in extracellular recordings in the ferret [8]. Those data are typically characterized by simple quantifications such as the firing rate or the mean subthreshold membrane potential [5], but a more comprehensive quantification is often missing. So the increasing amount of intracellular data of awake animals as well as the need to compare in a rigorous way the data under various recording conditions call for a model of spontaneous activity in single neurons.

Single neuron models have been studied for more than a century. Simple models such as the integrate-and-fire model [9, 10] and its more recent nonlinear versions [11–13] describe the relationship between the input current and the membrane potential in terms of a small number of parameters and are therefore convenient for analytical treatment, but do not provide much insight about the underlying biophysical processes. On the other end of the spectrum, biophysical models such as the Hodgkin-Huxley model [14, 15] relate the input current to the membrane potential through a detailed description of the various transmembrane ion channels, but estimating the model parameters remains challenging [16, 17]. Despite the success of those types of models, none of them can be directly applied to intracellular *in vivo* recordings for the simple reason that the input current is not known.

Another reason why a precise model of spontaneous activity is needed is that there are several theories that have been proposed that critically depend on statistical properties of spontaneous activity. For example Berkes et al. validate their Bayesian treatment of the visual system by comparing the spontaneous activity and the averaged evoked activity [8]. Another Bayesian theory proposed the idea that short-term plasticity acts as a Bayesian estimator of the presynaptic membrane potential [18]. To validate this theory, it is also necessary to characterize spontaneous activity with a statistical model that describes the subthreshold as well as the spiking dynamics.

The last motivation for a model that describes both the subthreshold and the suprathreshold dynamics is the possibility to separate those two dynamics in a principled way. Indeed, it is interesting to know from the recordings what reflects the input dynamics and what aspect comes from the neuron itself (or rather what is associated with the spiking dynamics). Of course a simple voltage threshold can separate the sub- and a suprathreshold dynamics, but the value of the threshold is somewhat arbitrary and could lead to undesirable artifacts. Therefore a computationally sound model that decides itself what belongs to the subthreshold and what belongs to the suprathreshold dynamics is highly desirable.

Here, we propose a single neuron model that describes intracellular *in vivo* recordings as a sum of a sub- and suprathreshold dynamics. This model is flexible enough in order to capture the large diversity of neuronal dynamics while remaining tractable, i.e. the model can be fitted to data in a reasonable time. More precisely, we propose a stochastic process where the subthreshold membrane potential follows a Gaussian process and the firing intensity is expressed as a non-linear function of the membrane potential. Since we further include refractoriness and adaptation mechanisms, our model, which we call the Adaptive Gaussian Point Emission process (AGAPE), can be seen as an extension of both the log Gaussian Cox process [19] and the generalized linear model [20–22].

Results

Here we present a statistical model of the subthreshold membrane potential and firing pattern of a single neuron *in vivo*. See Fig 1A for such an *in vivo* membrane potential recording. We first provide a formal definition of the model and then show a range of different results. 1) The model is flexible and supports arbitrary autocorrelation structures and adaptation kernels. Therefore, the range of possible statistical features is very large. 2) The model is efficiently fittable and the learning procedure is validated on synthetic data. 3) The model can be fitted to *in vivo* datasets. 4) All the features included in the model are required to provide a good description of *in vivo* data.

Definition of the AGAPE model

The AGAPE model is a single neuron model where the input to the neuron is not known, which is typically the case under *in vivo* conditions. The acronym AGAPE stands for Adaptive




Fig 1. (A) A sample *in vivo* membrane potential trace from an intracellular recording of a neuron in HVC of a Zebra Finch. (B) The generative AGAPE model can generate a trace of subthreshold membrane potential u (top trace). Based on this potential, a spike train s is generated (middle, dashed vertical lines). Finally, a stereotypic spike-related kernel is convolved with the spike train and added to u , giving rise to u_{som} (bottom, thick line). This quantity is the synthetic analog of the recorded, preprocessed *in vivo* membrane potential.

doi:10.1371/journal.pone.0142435.g001

GAussian Point Emission process since the subthreshold membrane potential follows a Gaussian process and since the spike emission process is adaptive.

More formally, the AGAPE process defines a probability distribution $p(u_{\text{som}}, s)$ over the somatic membrane voltage trace $u_{\text{som}}(t)$ and the spike train $s(t) = \sum_{i=1}^{n_s} \delta(t - \hat{t}_i)$ where \hat{t}_i , $i = 1, \dots, n_s$ are the nominal spike times (decision times), which occur a certain fixed time period $\delta > 0$ before the peak of the action potential. From this probability distribution (or generative model) it is possible to draw samples that look like intracellular *in vivo* activity (for practical purposes, the samples will be compared to the preprocessed recordings, see explanations below). The AGAPE model assumes that the somatic membrane voltage as a function of time $u_{\text{som}}(t)$ is given by (see Fig 1)

$$u_{\text{som}}(t) = u_r + u(t) + u_{\text{spike}}(t), \quad (1)$$

where u_r is a constant (the reference potential), $u(t)$ describes the subthreshold membrane potential as a stochastic function drawn from a stationary Gaussian process (GP) [23]

$$u \sim \mathcal{GP}[0, k(t - t')] \quad (2)$$

with covariance function $k(t - t')$ (which can be parametrized by a weighted sum of exponential decays with weights σ_i^2 and inverse time constants θ_i , see [Materials and Methods](#)). For small values of δ (e.g. 1–3 ms), $u(t)$ can be seen as the net contribution from the unobserved synaptic inputs and $u_{\text{spike}}(t)$ is the spike-related contribution (see Fig 1B) which consists of the causal convolution of a stereotypical spike-related kernel α with the spike train $s(t)$, i.e.

$$u_{\text{spike}}(t) = \int_0^\infty \alpha(t')s(t - t')dt'. \quad (3)$$

where α can be parametrized by a weighted sum of basis functions with weights a_i , see [Materials and Methods](#). Here, we have made a separation of subthreshold and suprathreshold layers, in that whatever is stereotypic and triggered by the point-like spikes $s(t)$ is attributed to $u_{\text{spike}}(t)$, and the rest belongs to the fluctuating signal $u(t)$. This separation need not correspond to the biophysical distinction between synaptic inputs and active processes of the recorded cell (i.e. the positive feedback loop of the spiking mechanism). Indeed, especially for a choice of large δ (e.g. ~ 20 ms), $u_{\text{spike}}(t)$ also contains large depolarizations due to strong synaptic input which cannot be explained by the GP signal $u(t)$.

Note that this model could easily be extended by including an additional term in Eq (1) which depends on an external input, e.g. a linear filter of the input (see also Discussion). However, since this input current was not accessible in our recordings, its contribution was assumed to be part of $u(t)$ or $u_{\text{spike}}(t)$.

Now we proceed to the coupling between the subthreshold potential $u(t)$ and the spiking output, as well as adaptive effects associated with spike generation. These effects are summarized by an instantaneous firing rate $r(t)$ —as in the generalized linear model (GLM) [20–22] or escape-rate models [24]—which is computed from the value of the subthreshold membrane potential at time t , $u(t)$, and the spike history as

$$r(t) = g[A(t) + \beta u(t)], \quad A(t) = \int_0^\infty \eta(t')s(t-t')dt', \quad (4)$$

where $\beta \geq 0$ is the coupling strength between u and the spikes, and $A(t)$ is the adaptation variable which is the convolution of an adaptation kernel η (which can be parametrized by a weighted sum of basis functions with weights w_i , see Materials and Methods) with the past spike train. Also note that we choose not to model adaptation currents explicitly, since they would simultaneously impact the membrane potential and the firing probability (see Discussion). The function g is called gain function, and here we use an exponential one, i.e. $g[A(t) + \beta u(t)] = e^{\log r_0 + A(t) + \beta u(t)}$. Other functional forms such as rectified linear or sigmoidal could be used depending on the structure of the data. However, this choice has important implications on the efficiency of learning of the model parameters [25]. We define the probability density for s on an interval $[0, T]$ conditioned on u as

$$p(s|u) = \exp\left(-\int_0^T r(t) dt\right) \prod_{i=1}^{n_s} r(\hat{t}_i), \quad s(t) = \sum_{i=1}^{n_s} \delta(t - \hat{t}_i), \quad 0 \leq \hat{t}_1 < \dots < \hat{t}_{n_s} \leq T, \quad n_s \in \mathbb{N}_0. \quad (5)$$

The parameter β connects the subthreshold membrane potential u to the rate fluctuations. The magnitude of the rate fluctuations depend on the variance σ^2 of u , and therefore we use $\beta\sigma$ as a measure of the effective coupling strength. When $\beta > 0$ the quantity $\theta(t) = -A(t)/\beta$ can be regarded as a soft threshold variable which is modulated after a spike, and $u(t) - \theta(t)$ is the effective membrane potential relevant for the spike generation. This spiking process is a point process which generalizes the log Gaussian Cox process. Indeed, when $A = 0$, Eq (5) describes an inhomogeneous Poisson process with rate $g[\beta u(t)]$.

Practically, if we want to draw a sample from the AGAPE process, we first draw a sample u from the Gaussian Process (see S1 Text for how to do this efficiently), then for each time t we draw spikes $s(t)$ with probability density $r(t)$ and update the adaptation variable $A(t)$. Finally, the somatic membrane potential is calculated using Eq 1.

It is important to emphasize at this point that while the model may be directly fitted to the raw membrane potential u_{raw} as recorded by an intracellular electrode, we median filter the data in order to avoid artifacts and downsample for computational efficiency (see ‘Materials and Methods’). In this study the model is always fitted to the preprocessed recordings u_{som}^* and this is reflected e.g. in the shape of α which is most strongly affected by the preprocessing. It is important to keep in mind this point while interpreting the results of model fitting. The details of the preprocessing steps which were used are given in the ‘Materials and Methods’ section.

The model has a rich dynamical repertoire

The AGAPE provides a flexible framework which can be adjusted in complexity to model a wide range of dynamics. While for the datasets presented here a covariance function was used which consists of a sum of Ornstein-Uhlenbeck (OU) kernels, the Gaussian Process (GP)

allows for arbitrary covariance functions to be used. This includes simple exponential decay (as produced by a leaky integrate-and-fire neuron driven by white noise current), but it can produce also more interesting covariance functions such as power-law covariances, which are reported in [7, 26], or subthreshold oscillations, as reported in [27].

The model is also able to reproduce a wide range of firing statistics. A common measure of firing irregularity is the coefficient of variation (C_V , i.e. the ratio of standard deviation and mean) of the inter-spike interval distribution. In the absence of adaptation, the AGAPE is a Cox process and therefore has a coefficient of variation $C_V \geq 1$ [28]. The precise value of the C_V is a function of the coupling strength ($\beta\sigma$) as well as the autocorrelation of the GP. To illustrate this, we sampled synthetic data from a simple version of the AGAPE where the subthreshold potential u is an OU process with time-constant τ . As shown in Fig 2A, the C_V is an increasing function of the membrane time-constant τ , baseline firing rate r_0 , and dimensionless coupling parameter between membrane potential and firing rate $\beta\sigma$. Moreover, the range of the C_V extends from 1 to ≈ 8 within a range of $\beta\sigma \in [0, 2]$ and $r_0\tau \in [2^{-2}, 2^8]$. In the presence of adaptation, firing statistics are markedly different and can produce values of $C_V < 1$ [24, 29]. To illustrate this point, we considered an exponential adaptation kernel, i.e. $\eta(t) = -\eta_0 e^{-t/\tau_r}$. While the C_V increases as a function of $\beta\sigma$ and $r_0\tau$ as before, the range of values of the C_V now also covers the interval (0, 1) which is not accessible by the Cox process but which is observed in many neurons across the brain [30]. In order to study the influence of the parameters of the adaptation mechanism, we fix $\beta\sigma = r_0\tau = 1$ and plot the C_V as a function of $r_0\tau_r$ and η_0 (see Fig 2B). Within the parameter region explored in Fig 2B, the C_V spans values from 0.1 up to 1.6.

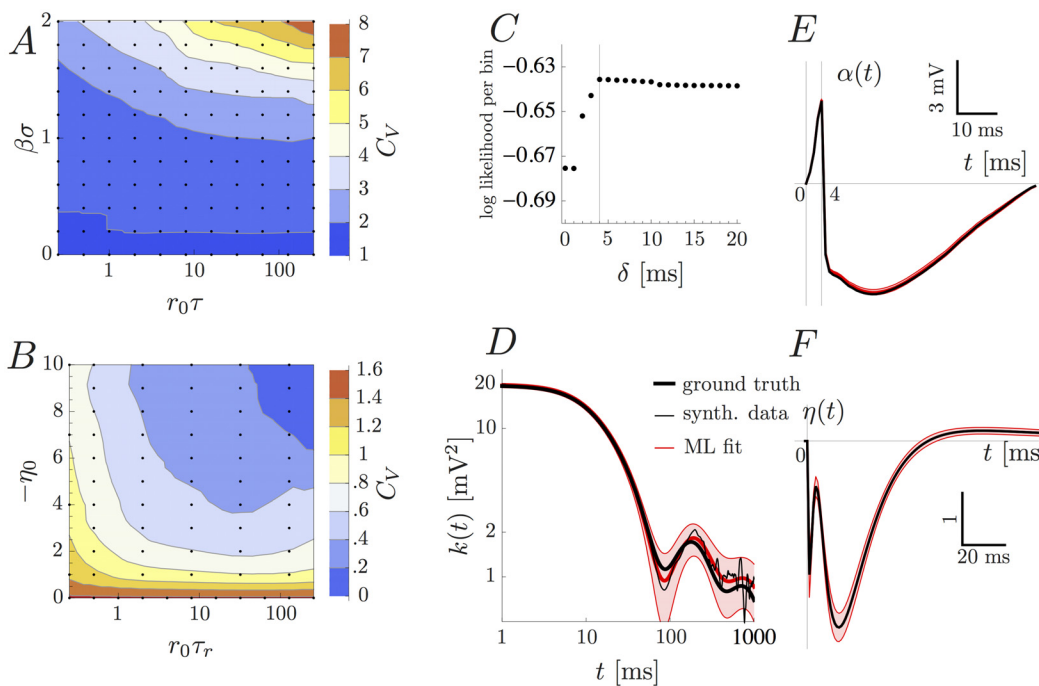


Fig 2. The model has a rich dynamical repertoire (A, B) and can be correctly fitted to synthetic data (C-F). (A, B) The coefficient of variation (C_V) of the inter-spike interval distribution is computed for parameter values shown as black dots and then linearly interpolated. (A) The C_V of a simple version of the AGAPE ($k(t) = \sigma^2 e^{-t/\tau}$, $\alpha = \eta = 0$) as a function of the model parameters (membrane time-constant τ , baseline firing rate r_0 and coupling strength $\beta\sigma$). (B) C_V of the AGAPE model with an exponentially adaptive process with fixed membrane time-constant, firing rate and coupling ($\beta\sigma = r_0\tau = 1$) as a function of the parameters describing adaptation, namely adaptation strength η_0 and time-constant τ_r . (C, D, E, F) Synthetic data is sampled from the AGAPE model with GP (D), spike-related (E), and adaptation (F) kernels as depicted in black, and $\delta = 4$ ms, $r_0 = 4.15$ Hz, $\beta = 0.374$ mV $^{-1}$. Then the AGAPE is fitted to the synthetic data by maximum likelihood (ML). (C) The maximum log likelihood per bin as a function of the parameter δ has its maximum at the ground truth value $\delta = 4$ ms. (D, E, F) The ML estimates (red) of the GP, spike-related and adaptation kernels lie within two standard deviations (red shaded regions, estimated by means of the observed Fisher information) from the ground truth.

doi:10.1371/journal.pone.0142435.g002

The model can be fitted efficiently

The parameters of the AGAPE model are learned through a maximum likelihood approach. More precisely, we fit the model to an *in vivo* sample (highlighted by a ^{*}) of preprocessed somatic membrane potential u_{som}^* and spike train $s^{*,\delta}$ by maximizing the log likelihood applied to the joint data set $(u_{\text{som}}^*, s^{*,\delta})$ over the parameter space of the model (i.e. $u_r, \log r_0, \beta$, the coefficients of the kernels k, η , and α , and the delay parameter δ). The empirical spike train $s^{*,\delta}$ depends on the parameter δ because the formal spike times \hat{t}_i are assigned to be a time period δ before the recorded peak of the action potential. The joint probability of the data can be expressed as a product

$$\begin{aligned}
 p(u_{\text{som}}^*, s^{*,\delta}) &= \int p(u)p(s^{*,\delta}|u)p(u_{\text{som}}^*|u, s^{*,\delta})\mathcal{D}u \\
 &= \int p(u)p(s^{*,\delta}|u)\delta(u_{\text{som}}^* - u - u_r - u_{\text{spike}})\mathcal{D}u \\
 &= p(u = u_{\text{som}}^* - u_r - u_{\text{spike}})p(s^{*,\delta}|u = u_{\text{som}}^* - u_r - u_{\text{spike}}) \\
 &\equiv p_{s^{*,\delta}}(u_{\text{som}}^*)p(s^{*,\delta}|u_{\text{som}}^*).
 \end{aligned}
 \tag{6}$$

The subscript $s^{*,\delta}$ of the first factor denotes the explicit dependence on the spike train. The individual terms on the r.h.s. will be given below. The function we are optimizing is the logarithm of the above joint probability which we can write as

$$\mathcal{L}(u_r, k, \alpha, \log r_0, \beta, \eta, \delta) = \log p_{s^{*,\delta}}(u_{\text{som}}^*; u_r, k, \alpha) + \log p(s^{*,\delta}|u_{\text{som}}^*; u_r, \alpha, \log r_0, \beta, \eta). \tag{7}$$

It should be noted that the presence of the spike-related kernel α in both terms produces a trade-off situation: removing the spike-related trajectory improves the Gaussianity of the membrane potential u (and therefore boosts the first term) at the cost of the of the second term by removing the short upward fluctuation that leads to the spike. This trade-off situation makes maximum likelihood parameter estimation a non-concave optimization problem. Moreover, the evaluation of the GP likelihood of n samples, where $n = \mathcal{O}(10^5)$, comes at a high computational cost. Two important techniques make the parameter learning both tractable and fast: the first is the use of the circulant approximation of the GP covariance matrix which makes the evaluation of the likelihood function fast. The second is the use of an alternating fitting algorithm which (under an appropriate parametrization, see ‘Materials and Methods’) replaces the non-concave optimization in the full parameter space with two concave optimizations and a non-concave one in suitable parameter subspaces. Those two techniques are further described in the next section.

Efficient likelihood computation. The log-likelihood function is evaluated in its discrete-time form with n time points separated by a time-step Δt . The GP variable u (which leads to u_{som} through Eq (1)) is multivariate Gaussian distributed with a covariance matrix $K_{ij} = k(t_i - t_j)$, where $t_i = i\Delta t$. The matrix K is symmetric and, by virtue of stationarity, Toeplitz. Evaluation of the GP likelihood requires inversion of K , which is computationally expensive (the time required to invert a matrix typically scales with n^3). For this reason we approximate this Toeplitz matrix by the circulant matrix C which minimizes the Kullback-Leibler divergence (see [31–33] and S1 Text)

$$C = \underset{D \text{ circulant}}{\operatorname{argmin}} \mathcal{D}_{\text{KL}}[\mathcal{N}(m, K)||\mathcal{N}(m, D)] \tag{8}$$

between the two multivariate Gaussian distributions with the same mean but different

covariance matrices. This optimization problem can be solved by calculating the derivative of $\mathcal{D}_{\text{KL}}[\mathcal{N}(m, K) || \mathcal{N}(m, D)]$ with respect to D and using the diagonalization of D by a Fourier transform matrix [33]. After a bit of algebra (see S1 Text), denoting $k_i = K_{1i}$ and $k_{n+1} \equiv 0$, the optimal circulant matrix can be written as $C_{ij} = c_{(i-j \bmod n)+1}$, where $i, j = 1, \dots, n$ and

$$c_i = \frac{1}{n} [(n - i + 1)k_i + (i - 1)k_{n-i+2}]. \tag{9}$$

The replacement of K by C is equivalent to having periodic boundary conditions on u , which has a small effect under the assumption that the time interval spanned by the data is much longer than the largest temporal autocorrelation length of k . So the first term on the r.h.s. of Eq (6) is a multivariate Gaussian density $\mathcal{N}(0, C)$. The determinant of the covariance matrix C is the product of eigenvalues, which for a circulant matrix are conveniently given by the entries of \hat{c} , the discrete Fourier transform of c (see the S1 Text for our conventions regarding discrete Fourier transforms). Also the scalar product $u^T C^{-1} u$ can be written in terms of \hat{c} . Together, the first term on the r.h.s. of Eq (6) takes the simple form

$$\log p_{s^*,\delta}(u_{\text{som}}^*) = -\frac{1}{2} \sum_{i=1}^n \left(\log(2\pi\hat{c}_i) + \frac{1}{n} \frac{|\hat{u}_i|^2}{\hat{c}_i} \right), \tag{10}$$

where \hat{u}_i are the components of the discrete Fourier transform of u^* . The Gaussian component of the membrane potential u is implicitly given by the discretized somatic voltage modified by a discrete-time version of the spike-related kernel convolution,

$$u_i^* = u_{\text{som},i}^* - u_r - \sum_{j=1}^{i-1} \alpha_j s_{i-j}^{*,\delta}, \tag{11}$$

where $s_i^{*,\delta}$ is the binned spike train (see below), α_i is a discretized version of the spike-related kernel. The time required to compute $\log p_s(u_{\text{som}}^*)$ is determined by the complexity of the Fourier transform, which is of the order of $n \log n$. This dramatic reduction in complexity (compared to n^3) allows a fast evaluation of the log-likelihood.

The spiking distribution $p(s^{*,\delta} | u_{\text{som}}^*)$ is approximated by a Poisson distribution with constant rate within one time bin. For each bin, $s_i^{*,\delta}$ counts the number of spikes that occur in that bin, and the conditional likelihood of the spikes therefore reads

$$\log p(s^{*,\delta} | u_{\text{som}}^*) = \sum_{i=1}^n \{ s_i^{*,\delta} \log [r_i \Delta t] - r_i \Delta t - \log [s_i^{*,\delta}!] \}, \tag{12}$$

where $r_i = g[\beta u_i^* + \sum_{j=1}^{i-1} \eta_j s_{i-j}^{*,\delta}]$ and u_i^* as defined in Eq (11). If $s_i^{*,\delta}$ contains only zeros and ones (which can be accomplished given small enough bins), the last term $\log s_i^{*,\delta}!$ vanishes.

Efficient parameter estimation. Except for the parameter δ , which takes discrete values of multiples of the discretization step Δt , it is possible to analytically calculate the first and second partial derivatives of the likelihood function defined in Eq (6) with respect to the model parameters $(u_r, k, \alpha, \log r_0, \beta, \eta)$ (see S1 Text) to facilitate the use of gradient ascent, Newton, and pseudo-Newton optimization algorithms. A desirable feature of an optimization problem is concavity of the objective function (in our case, the log-likelihood function). Even though the problem of finding optimal parameters for the AGAPE process is not concave, the optimization can be done in three alternating subspaces (see Fig 3). The full set of parameters Θ is divided into three parts: θ_{GP} for the GP parameters (u_r , parameters of k), $\theta_{\text{spike kernel}}$ for the spike-related kernel parameters, and θ_{spiking} for the parameters controlling spike emission ($\log r_0, \beta$ and parameters of η). The optimization then proceeds according to the following cycle: (1)

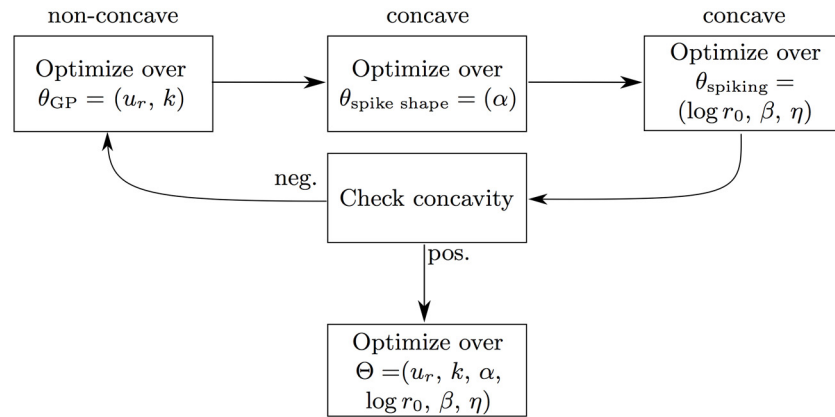


Fig 3. This schematic shows the optimization scheme that is used to learn the parameters of the AGAPE model when it is fitted to the data (for a given δ). As long as the current parameter estimate sits in a non-concave region of the likelihood function, the top cycle optimizes over different subspaces of the parameter space. If and when a concave point is reached, the optimization proceeds in the full parameter space. This whole scheme is repeated for each value of δ in order to find the optimal one.

doi:10.1371/journal.pone.0142435.g003

the GP parameters are learned, (2) the spike-related kernel parameters are learned, and lastly (3) the spiking parameters are learned. In each step the remaining parameters are held fixed. The cycle is repeated until the parameters reach a region where the log likelihood is locally concave in the full parameter space, after which the optimization can be run in the full parameter space until it converges. Joint concavity of the log likelihood holds if all the eigenvalues of the Hessian matrix are strictly negative. As shown in [25], step (3) is concave for a certain class of gain functions g , including the exponential function, and linear parametrizations of the adaptation kernel. The same holds for the spiking term of the log-likelihood in step (2). The voltage term of the log likelihood of step (2) is concave by numerical inspection in the cases we considered. To summarize, steps (2) and (3) are concave and Newton’s method can be used in these steps as well as for the final concave optimization in the full space. Step (1) is non-concave and therefore a simple gradient ascent algorithm is used.

The optimization over $(u_r, k, \alpha, \log r_0, \beta, \eta)$ is repeated for every $\delta = 0, \Delta t, 2\Delta t, \dots, \delta_{\max}$ in order to select the one δ that maximizes the log-likelihood $\mathcal{L}(u_r, k, \alpha, \log r_0, \beta, \eta, \delta)$. The value of δ_{\max} is chosen such that it is less than the least upper bound of the support of the basis of the spike-related kernel α . Since the parameters $u_r, k, \alpha, \log r_0, \beta, \eta$ are expected to change only a little when going from one δ to the next, $\delta + \Delta t$, learned parameters for δ can be used as initial guesses for nearby $\delta + \Delta t$ or $\delta - \Delta t$. We thus get two different initializations, which we can exploit by starting e.g. with $\delta = 0$, ascending through the sequence of candidate δ ’s up to the maximum δ , and descending back to zero.

Validation with synthetic data

Despite this improvement in speed and tractability, the optimization is still riddled with multiple local minima which require the use of multiple random initializations. In order to demonstrate the validity of the fitting method, synthetic data of length 270.112 seconds ($n = 270112$, the same as *in vivo* dataset \mathcal{D}_1 , see below) was generated with known parameters ($\delta = 4$ ms, $r_0 = 4.15$ Hz, $\beta = 0.374$ mV⁻¹ and GP, spike-related kernel and adaptation kernels as depicted in Fig 2D–2F). The learning algorithm was initialized with least-squares estimates of the covariance function parameters σ_i^2 based on the empirical autocorrelation function of u_{som} and spike-related kernel and adaptation kernels set to zero. The true underlying δ can be recovered

from the synthetic data (Fig 2C). Moreover, the algorithm converges after a few dozen iterations (taking only three minutes on an ordinary portable computer) and—with δ set to 4 ms—recovers the correct GP, spike-related, and adaptation kernels (Fig 2D–2F). All ML estimates lie within a region of two standard deviations around the ground truth, where standard deviations are estimated from the observed Fisher information [34].

The model can fit *in vivo* data

We fitted the model to a number of *in vivo* traces from different animals and conditions (see ‘Materials and Methods’ for a detailed description of the data sets). We would like to remind the reader at this point that the model is never fitted to the raw membrane potential, but to a preprocessed, i.e. median-filtered and downsampled dataset (see Materials and Methods). Because of this preprocessing stage, the model only sees the truncated action potentials which emerge from the median filter. This is reflected in the extracted spike-related kernel α , which is characterized by a smaller amplitude than the original action potential in the raw membrane potential data.

We show the detailed results of the model fitting for the example songbird HVC dataset \mathcal{D}_1 . The optimal value of δ for this dataset was $\delta = 18$ ms (see S3 Fig), with which the model captures the subthreshold and suprathreshold statistics (smaller values of δ compromise both the subthreshold and suprathreshold description because the large upward fluctuations which precede spikes in this dataset are unlikely to arise from a GP). In particular, the stationary distribution of the membrane potential u is well approximated by a Gaussian (Fig 4B) and pronounced after-hyperpolarization is seen in the spike-related kernel (Fig 4D). The subthreshold autocorrelation structure is well reproduced by the parametric autocorrelation function k (Fig 4C). The adaptation kernel reveals an interesting structure in the way the spiking statistics deviates from a Poisson process (Fig 4E). This feature of the spiking statistics is also reflected in the inter-spike interval (ISI) distribution (Fig 4F). Both the data and the fitted model first show an increased, and then a significantly decreased probability density when compared to a pure Poisson process. The remaining parameters are listed in Table 1 (errors denote two standard deviations, estimated from Fisher information, see Materials and Methods). The model can be used to generate synthetic data, which is shown in Fig 4H.

In order to show the generality of the model, we fitted the model on two more datasets, \mathcal{D}_3 from another HVC neuron and \mathcal{D}_4 from mouse visual cortex. The parameter δ was found to take the optimal value of 12 ms for \mathcal{D}_3 and 32 ms for \mathcal{D}_4 (to see how fitted parameters change as a function of δ , see S3 and S4 Figs). The comparison of the GP, spike-related and adaptation kernels is shown in Fig 5, and the remaining parameters are listed in Table 1. The three cells show pronounced differences in autocorrelation structure, spike-related kernel and spike-history effects. In particular the two datasets \mathcal{D}_1 and \mathcal{D}_4 show rather long autocorrelation lengths of the membrane potential and asymmetric spike-related kernels, whereas the cell in \mathcal{D}_3 has comparatively short autocorrelation length and very pronounced hyperpolarization. Adaptation is much stronger in \mathcal{D}_3 than in \mathcal{D}_1 , balancing the much higher baseline firing rate r_0 , see Table 1. The error bars on the adaptation kernel are small for datasets \mathcal{D}_1 and \mathcal{D}_3 due to the abundance of spikes. On the other hand, the adaptation kernel of dataset \mathcal{D}_4 is poorly constrained by the available data. This is due to the fact that dataset \mathcal{D}_4 consists of very short trials with very few spikes. Despite this fact, good agreement is achieved between the distribution of inter-spike intervals of the *in vivo* data and ISI statistics sampled from the AGAPE (see Fig 5, bottom row) for all datasets.

The model does not overfit *in vivo* data

The AGAPE process has a fairly large number of parameters. Therefore it is important to check whether the model overfits the data, compromising its generalization performance. In

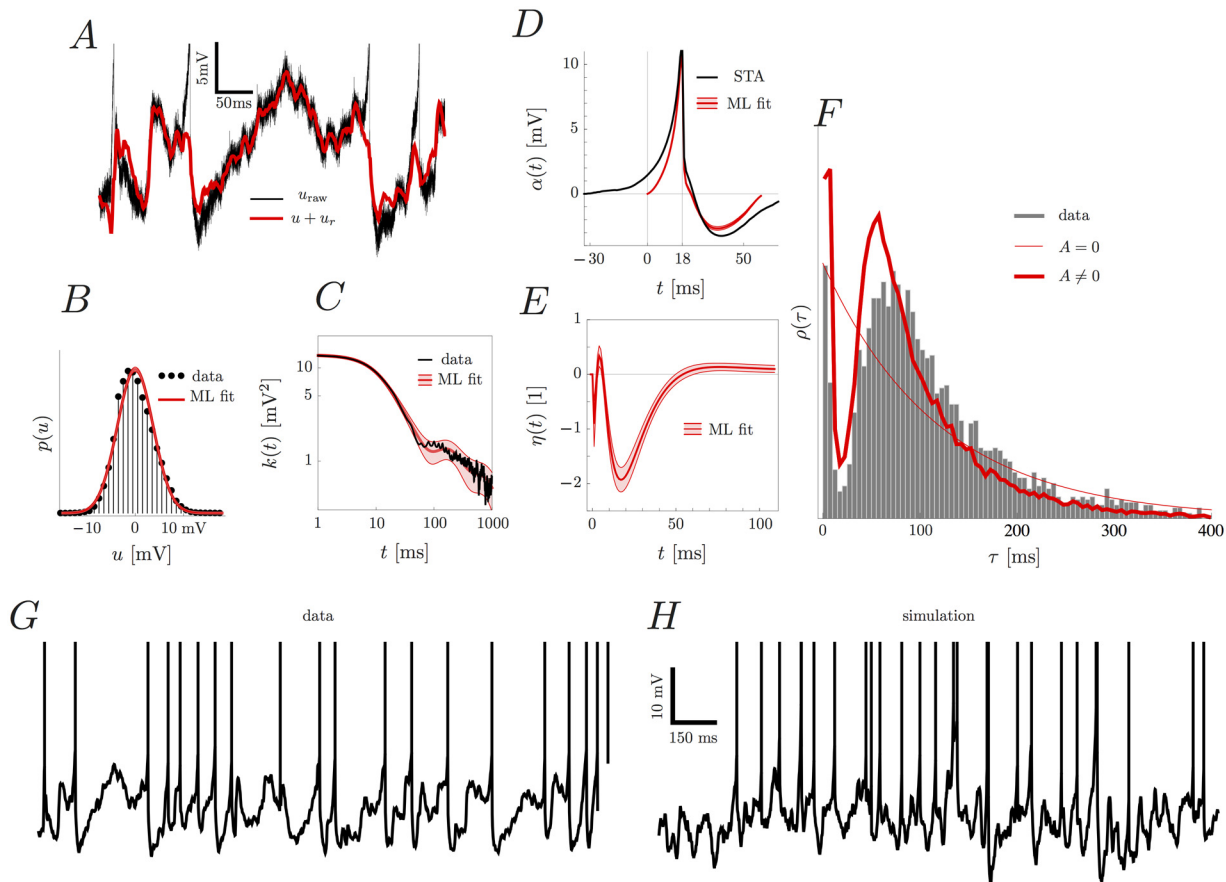


Fig 4. The results of maximum likelihood (ML) parameter fitting to dataset \mathcal{D}_1 . After fitting, we see (A) the removal of the spike-related kernel through the difference between the recorded trace u_{raw} and the subthreshold membrane potential $u+u_r$; (B) the match of the stationary distribution of the subthreshold potential u and a Gaussian. We also observe that (C) the autocorrelation function of the data, Eq (14), is well reproduced by $k(t)$ in Eq (15); (D) the spike-related kernel $\alpha(t)$ starts at $-\delta = -18$ ms relative to the peak of the action potential. The difference between the spike-triggered average (STA) and the spike-related kernel is attributed to the GP; and (E) that the adaptation kernel $\eta(t)$ shows distinct modulation of firing rate which produces firing statistics significantly different from a Poisson process. This is also reflected in the inter-spike interval density $\rho(\tau)$ (F) of the data, which shows good qualitative agreement with a simulated AGAPE with adaptive kernel as in (E) (thick red line), but not by a non-adaptive (i.e. Poisson) process (thin red line). After fitting, a two second sample of synthetic data (H) looks similar as the *in vivo* data (G). In (G, H) vertical lines are drawn at the spiking times. All red shaded regions denote ± 2 standard deviations, estimated from the observed Fisher information.

doi:10.1371/journal.pone.0142435.g004

Table 1. The values (p.m. two standard deviations, estimated from the observed Fisher information) of the fitted parameters not shown in Fig 5 for the *in vivo* datasets described in the main text. The last row shows the effective coupling strength between the membrane potential and the firing rate, given by β times the standard deviation σ of the membrane potential.

Dataset	\mathcal{D}_1	\mathcal{D}_3	\mathcal{D}_4
δ [ms]	18	12	32
u_r [mV]	-52.9 ± 0.2	-66.6 ± 0.1	-51.5 ± 0.5
r_0 [Hz]	11.7 ± 0.6	71 ± 5	0.15 ± 0.05
β [mV^{-1}]	0.12 ± 0.01	0.24 ± 0.02	0.46 ± 0.05
$\beta\sigma$ [1]	0.45 ± 0.03	0.67 ± 0.04	1.3 ± 0.1

doi:10.1371/journal.pone.0142435.t001

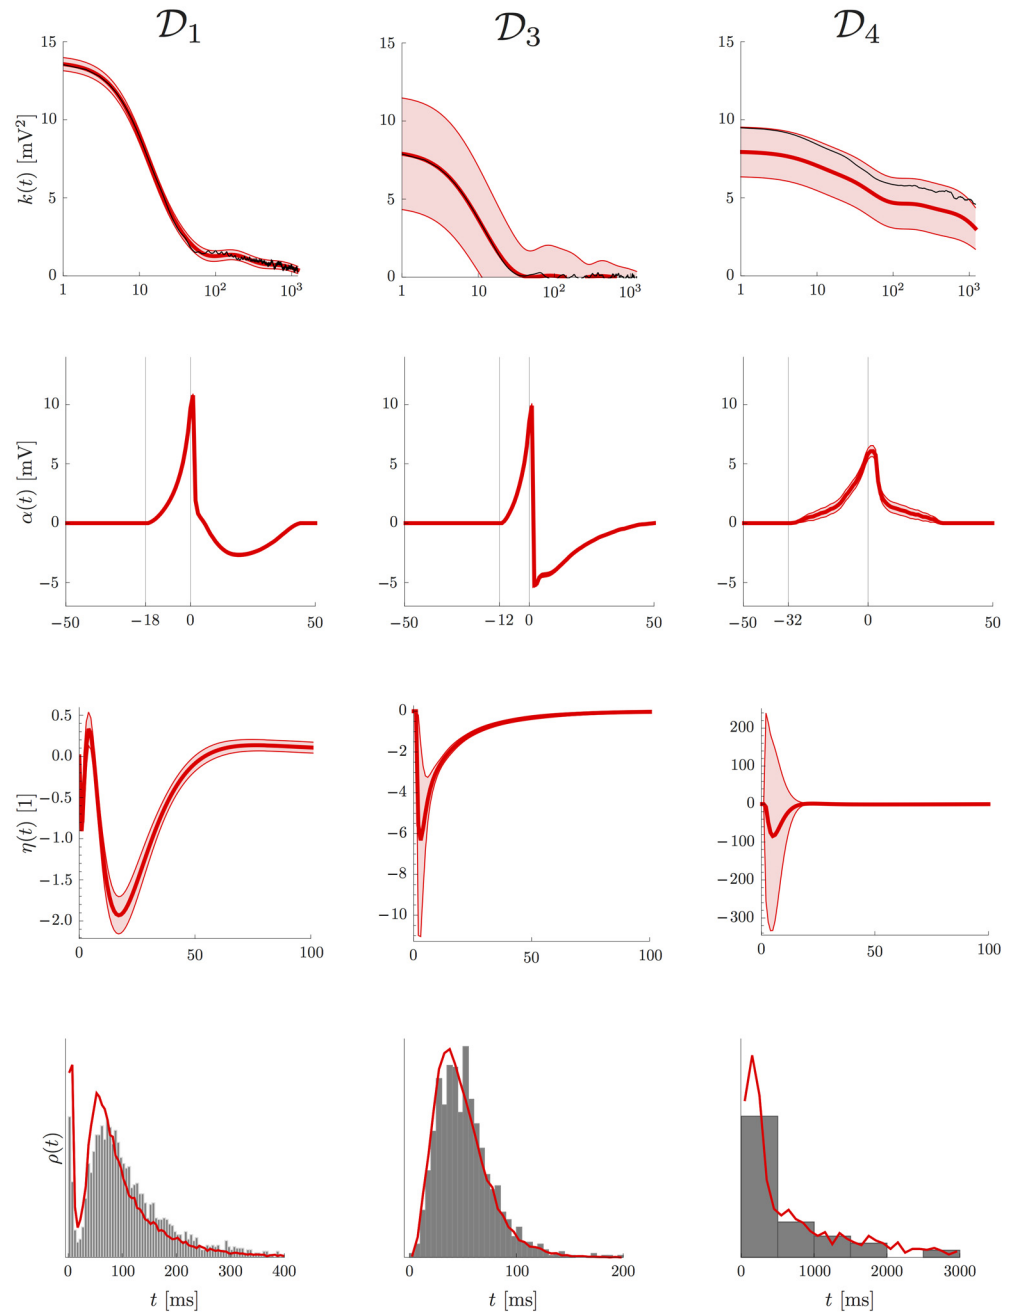


Fig 5. Fitting results for three different datasets. Dataset \mathcal{D}_1 is the same as in Fig 4, i.e. an HVC neuron from anesthetized Zebra Finch. \mathcal{D}_3 is from HVC in awake Zebra Finch, and \mathcal{D}_4 is from mouse visual cortex in awake mouse. The different panels show the results after fitting; in the first line the GP covariance function $k(t)$ (red) and the empirical autocorrelation (black), Eq (14), in the second line the spike-related kernel $\alpha(t)$, in the third line the adaptation kernel $\eta(t)$, and in the fourth line the inter-spike interval density $\rho(t)$ (data ISI histogram in gray, simulated ISI distribution from AGAPE in red). There are pronounced differences between datasets in all three kernels, showing the flexibility of the AGAPE model in describing a wide range of statistics. All red shaded regions denote ± 2 standard deviations, estimated from the observed Fisher information.

doi:10.1371/journal.pone.0142435.g005

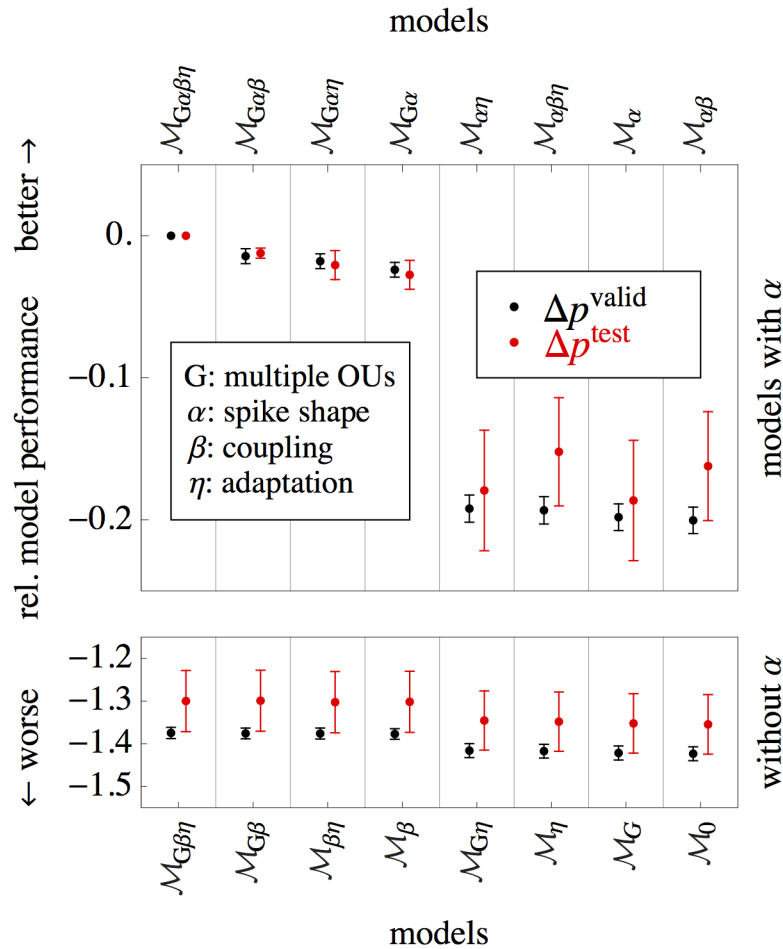


Fig 6. Comparison of the different models on dataset \mathcal{D}_2 . The relative measure of model performance, i.e. the per-bin log-likelihood Δp (see Eq (13)) between any model and the most complex model ($\mathcal{M}_{G\alpha\beta\eta}$) are significantly negative (with exception of $\mathcal{M}_{G\alpha\eta}$, and trivially $\mathcal{M}_{G\alpha\beta\eta}$), implying that the added complexity improves the model fit without overfitting. This holds for both validation scores Δp^{valid} (black) and scores from unseen test data Δp^{test} (red). Error bars denote one standard error of the mean (S.E.M.). The biggest improvement of fit quality is achieved by including the spike-related kernel (upper vs. lower part of the figure).

doi:10.1371/journal.pone.0142435.g006

short, when a model has too many parameters, it tends to be poorly constrained by the data and therefore when the model is first trained on one part of the data and then tested on another part on which it is not trained, the test performance will be significantly worse than the training performance.

Here, we use cross-validation to perform a factorial model comparison on an exemplary dataset in order to validate the different structural parts of the model. The procedure is described in detail in the Materials and Methods.

Model comparison is performed on the dataset \mathcal{D}_2 and the results are shown in Fig 6, where the mean difference of per-bin log-likelihood (see ‘Materials and Methods’)

$$\Delta p_i^{\text{valid}} = \langle \Delta p_{ij}^{\text{valid}} \rangle_j, \quad \Delta p_i^{\text{test}} = \langle \Delta p_{ij}^{\text{test}} \rangle_j, \quad \Delta p_{ij}^{\text{valid,test}} = p_{ij}^{\text{valid,test}} - p_{G\alpha\beta\eta,j}^{\text{valid,test}} \quad (13)$$

is shown for all models $i \in \{0, \dots, G\alpha\beta\eta\}$ (here, $\langle \cdot \rangle_j$ denote averages over chunks j of the cross-validation). The results are very similar for both validation data (which was left-out during training, but appeared in other training runs) and the test data which was never seen during

training. The most complex model ($\mathcal{M}_{G\alpha\beta\eta}$) performs significantly better than any one of the simpler models on validation data except $\mathcal{M}_{G\alpha\eta}$ where the difference is too small and lies inside a region of two standard errors of the mean. This confirms that most of the model features are required to provide an accurate description of the experimental data.

Discussion

In this study, we introduced the AGAPE generative model for single-neuron statistics in order to describe the spontaneous dynamics of the somatic potential without reference to an input current. We showed that this model has a rich dynamical repertoire and can be fitted to data efficiently. By fitting a heterogeneous set of data, we finally demonstrated that the AGAPE model can be used for the systematic characterization and comparison of *in vivo* intracellular recordings.

Flexibility and tractability of the model

The AGAPE model provides a unified description of intracellular dynamics, offering a large degree of flexibility in accounting for the distinct statistical features of a neuron. As the example datasets demonstrate, the model readily teases apart the differences in the statistics which exist between different cells in different animals (see Fig 5). This shows that the model is sensitive enough to distinguish between datasets which are in fact very similar.

We used a set of approximations and techniques to make the model fitting tractable, despite the non-concavity of the log likelihood function. It is still the case that multiple local maxima of the likelihood function can make the fitting somewhat hard, especially if the quantity of data available for fitting is quite low. However, since one run of the fitting itself takes only a few minutes even on a portable computer, multiple initializations can be tried out in a relatively short amount of time.

Comparison with existing models

From an operational perspective, existing spiking neuron models can be divided into three main categories: stimulus-driven, current-driven and input-free spiking neurons. The first category contains phenomenological models that relate sensory stimuli to the spiking output of the neuron. The linear-nonlinear-Poisson model (LNP) [35], the generalized linear model (GLM) [20–22] or the GLM with additional latent variables [36] are typical examples in this category. Even though the spike generation of the AGAPE shares some similarities with those models, there is an important distinction to make. In those models the convolved input (i.e. the output of the ‘L’ step of the LNP or the input filter of the GLM) is an internal variable that does not need to be mapped to the somatic membrane potential whereas in our case, the detailed modeling of the membrane potential dynamics is an important part of the AGAPE. Consequently, those phenomenological models are descriptions of extracellular spiking recordings whereas the AGAPE models the dynamics of the full membrane potential accessible with intracellular methods.

The second class of spiking models aims at bridging the gap between the input current and the spiking output. The rather simple integrate-and-fire types of models such as the exponential integrate-and-fire [13] or the spike-response model [24, 37] as well as the more biophysical models such as the Hodgkin-Huxley model [15] fall within this category. In contrast to those models where the action potentials are caused by the input current, the AGAPE produces a fluctuating membrane potential and stochastic spikes without a reference to an input current.

The last category of models aims at producing spontaneous spiking activity without an explicit dependence to a given input [18, 38, 39]. For example, Cunningham et al. propose a doubly stochastic process where the spiking generation is given by a gamma interval process and the firing intensity by a rectified Gaussian process, which provides a flexible description of the firing statistics [38]. However, the membrane potential dynamics is not modeled. In opposition, the neuronal dynamics assumed by Pfister et al. [18] models explicitly the membrane potential (as a simple Ornstein-Uhlenbeck process) but is not flexible enough to capture the dynamics of *in vivo* recordings. Also any of the current-driven spiking neuron models mentioned above can be turned into an input-independent model by assuming some additional input noise. So why is there a need to go beyond stochastic versions of those models? An integrate-and-fire model with additive Gaussian white noise is certainly fittable, but does not have the flexibility to model arbitrary autocorrelation for the membrane potential. At the other end of the spectrum, a Hodgkin-Huxley model with some colored noise would certainly be able to model a richer dynamical repertoire, but the fitting of it remains challenging [16] (but see [17]). The main advantage of the AGAPE is that it is at the same time very flexible and easily fittable. The flexibility mostly comes from the fact that any covariance function can be assumed for the GP process. The relative ease of fitting comes from the circulant approximation as well as from the presence of concave subspaces in the full parameter space.

Another distinct feature of our model with respect to other existing models is the explicit modeling of the spike-related trajectory instead of the spike-triggered average (as e.g. in [40]). Even though both concepts share similarities—both would capture a sudden and strong input that lead to a spike—there is an important distinction. The spike-triggered average also captures the (possibly smaller) upward fluctuations of the membrane potential which causes the spike while the spike-related kernel α precisely avoids capturing those fluctuations, letting the GP kernel explain them.

So if we removed the spike-triggered average e.g. in synthetic data where the true coupling parameter β is large, we would also remove the characteristic upward fluctuation of the membrane potential which causes the spike. By doing so, the fitting procedure would not find the correct relation between the values of the membrane potential and the observed spike patterns and therefore choose a β close to zero. Thus, if something has to be removed around an action potential (and our model comparison, Fig 6, demonstrates convincingly that this is necessary), the formulation of the model demands that it is parametrically adjustable. This is the main reason why in our model framework the spike-triggered average has to be rejected as a viable extraction method. Note that if the true coupling parameter β is close to zero, the spike-triggered average is close to the extracted spike-related kernel α . For data where the action potential shape shows considerable variability, the model could be generalized to include a stochastic or a history-dependent spike-related kernel.

Extensions and future directions

Despite the focus of the present work on single-neuron spontaneous dynamics, the AGAPE model admits a straightforward inclusion of both stimulus-driven input and recurrent input. The inclusion of stimulus-driven input is similar as for the GLM model and allows the model to capture the neuronal correlate of stimulus-specific computation. The recurrent input makes the framework adaptable to multi-neuron recordings *in vivo*. While intracellular recordings from many neurons *in vivo* are very hard to perform, the rapid development of new recording techniques (e.g. voltage-sensitive dyes) makes the future availability of subthreshold data with sufficient time-resolution at least conceivable. The full-fledged model would allow questions regarding the relative importance of background activity, recurrent activity due to computation

in the circuit, and activity directly evoked by sensory stimuli to be answered in a systematic way. In this setup, the contribution of the GP-distributed membrane potential to the overall fluctuations would be reduced (since it has to capture less unrecorded neurons) while the contribution of the recorded neurons would increase. This modified model can be seen as a generalization of the stochastic spike-response model [24] or a generalization of the GLM (if the internal variable of the GLM is interpreted as the membrane potential).

So far, we assumed that weak synaptic inputs are captured by the Gaussian process while the strong inputs that lead to the spikes are captured by the spike-related kernel α . A straightforward extension of the model would be to consider additional intermediate inputs that cannot be captured by the GP nor by the spike-related kernel α but that can drive the neuron to emit (with a given probability) an action potential. Those intermediate input could be modeled as filtered Poisson events. The inclusion of those latent events would increase the complexity of the model and at the same time change some of the fitted parameters. In particular, we expect that it would increase the coupling β between the membrane potential and the firing rate and reduce the optimal delay δ between the decision time and the peak of the action potential. This could also provide a better way to separate the subthreshold dynamics (which depends on the input activity) from the suprathreshold dynamics (which would depend only on the neuron dynamics, and not on the strong inputs that it receives, as it is the case now).

A central assumption of our model is that of a Gaussian marginal distribution of the subthreshold potential. Although it is remarkably valid for the dataset considered here (i.e. the HVC dataset \mathcal{D}_1 see also Fig 4B), datasets characterized by a distinctly non-Gaussian voltage distribution even after spike-related kernel removal are beyond the scope of the current model. In order to address this limitation, the Gaussian process could be extended to a different stochastic process, e.g. a nonlinear diffusion process, permitting non-Gaussian and in fact arbitrary marginal distributions. Moreover, a reset behavior similar to the one exhibited by an integrate-and-fire model [13] could be achieved with a non-stationary GP which features a mean which is reset after a spike. Both modifications would have a severe impact on the technical difficulty of model fitting. Therefore, the Gaussian assumption can be regarded as a useful compromise which is preferable over a perfect account for the skewness of the marginal distribution.

The spike-related kernel method to separate subthreshold and suprathreshold dynamics is an important feature of the model which is used to rid the membrane potential recording of stereotypic waveforms associated with a spike. The spike related kernel as modeled in the AGAPE has no bearing on the probability of the spikes, whereas the adaptation kernel η which modulates the firing rate after a spike is not visible in the somatic membrane potential dynamics. A simple extension of the model could include spike-triggered adaptation currents which affect both the somatic membrane potential as well as the firing intensity. Another possible extension is to allow the firing probability to depend on a filtered version of the subthreshold potential u instead of the instantaneous value of u at a time δ before the peak of the action potential. Both of the mentioned extensions would improve the biophysical interpretability of the AGAPE, but they would also vastly increase the number of parameters. Therefore, a model comparison would be required to determine what level of model complexity is required in order to characterize the statistics of the recording.

In the present study, the AGAPE was fit to different datasets of two different animals and brain regions. A systematic fitting to *in vivo* intracellular data from a wide range of animals and brain regions would constitute a classification scheme which does not only complement existing classifications of neurons which are based on electrophysiological, morphological, histological, and biochemical data; such as the one in [41], but which is in direct relationship with the computational tasks the brain is facing *in vivo*.

Another application of the AGAPE could be in the context of a normative theory of short-term plasticity. Indeed, it has been recently hypothesized that short-term plasticity performs Bayesian inference of the presynaptic membrane potential based on the observed spike-timing [18, 42]. According to this theory, short-term plasticity properties have to match the *in vivo* statistics of the presynaptic neuron. Since the AGAPE provides a realistic generative model of presynaptic activity under which inference is supposedly performed, our model can be used to make testable predictions on the dynamical properties of downstream synapses.

Materials and Methods

Description of the datasets used

1. Dataset \mathcal{D}_1 is a recording from a HVC neuron of an anesthetized Zebra Finch (Ondracek and Hahnloser, unpublished recordings). The recording has a total length of 270 seconds at 32 kHz (see Fig 1A for a snippet of this recording) and contains 2281 action potentials.
2. Dataset \mathcal{D}_2 is another recording from a projection cell in HVC of Zebra Finch, but this time the animal is awake (Vallentin and Long, unpublished recordings). It consists of 6 individual recordings which together have a length of 152.5 seconds at 40 kHz. This dataset is used for model comparison (see below).
3. Dataset \mathcal{D}_3 is from similar conditions as \mathcal{D}_2 (Vallentin and Long, unpublished recordings, see [43, 44, 45] for similar recordings) and has a length of 60 seconds.
4. Dataset \mathcal{D}_4 consists of 19 individual trials of 4.95s duration at 20 kHz. The recording was obtained from a pyramidal neuron in layer 2/3 of awake mouse visual cortex [46].

Preprocessing

Intracellular voltage traces are often recorded at a rate between 20 and 40 kHz. This allows the action potentials to be resolved very clearly and precise spike timings to be extracted. However, for the study of the subthreshold regime, this high sampling rate is not required, and therefore the data may be down-sampled to roughly 1 kHz after obtaining the precise spike timings. Prior to down-sampling, we smooth with a median filter of 1ms width in order to truncate the sharp action potential peaks and avoid artifacts (see details below).

We define the spike peak times $\hat{t}_{i,\text{peak}}$ operationally as the time where the local maximum of the action potential is reached. This means that $\hat{t}_{i,\text{peak}}$ occurs after action potential onset, and hence the spike-related kernel has to extend to the past of $\hat{t}_{i,\text{peak}}$. The spike-related kernel starts at the nominal spike time \hat{t}_i which is shifted from the peak time by a fixed amount δ , i.e. $\hat{t}_{i,\text{peak}} = \hat{t}_i + \delta$. The nominal spike times \hat{t}_i are then binned to 1 ms, yielding a binary spike train $s_i = 0, 1$.

For $u_{\text{som}}(t)$ we use a preprocessed version of the recorded trace which has been median-filtered with a width of the filter of 1 ms and then down-sampled to 1 kHz, making it the same length as the binary spike train. This procedure preserves the relevant correlation structure of the membrane potential while reducing the computational demands of fitting as much as possible. In the data we examined, the median-filtered membrane potential has a dip after $\hat{t}_{i,\text{peak}}$ but unless down-sampling is done carefully, this dip sometimes occurs one timestep after $\hat{t}_{i,\text{peak}}$ and sometimes right at $\hat{t}_{i,\text{peak}}$ in the downsampled u_{som} . Since this dip will have to be captured by the spike-related kernel which has a fixed shape for all action potentials, the down-sampling

procedure has to ensure that the dip occurs always in the first time-step. We solved this problem by setting the down-sampled value of u_{som} at $\hat{t}_{i,\text{peak}}$ (rounded to 1 ms) to the value of u_{som} at $\hat{t}_{i,\text{peak}}$ before down-sampling.

While applying the model to the raw recording u_{raw} directly (without first filtering and downsampling it) is possible in principle, it comes at a massively increased computational cost. In the interest of time required to fit the model and amount of data having to be handled, it is therefore sensible to include that pre-processing stage.

Parametrizations and initializations

We already introduced the parameters u_r , r_0 and β . Additional parameters are needed to describe the autocorrelation $k(t)$, the spike-related kernel $\alpha(t)$ and the adaptation kernel $\eta(t)$.

The covariance function of the GP has to be parametrized such that it can explain the autocorrelation structure of the data. Therefore, an initial examination of the empirical autocovariance of u_{som} , i.e.

$$k^{\text{emp}}(j\Delta t) = \frac{1}{n-j-1} \sum_{i=1}^{n-j} \left(u_{\text{som},i} - \frac{1}{n-j} \sum_{k=1}^{n-j} u_{\text{som},k} \right) \left(u_{\text{som},i+j} - \frac{1}{n-j} \sum_{k=1}^{n-j} u_{\text{som},k+j} \right), \quad (14)$$

for $j = 0, \dots, j_{\text{max}}$, is done in order to determine a suitable basis. Here, we used a sum of Ornstein-Uhlenbeck (OU) kernels, i.e.

$$k(t) = \sum_{i=1}^{n_k} \sigma_i^2 e^{-\theta_i |t|}, \quad (15)$$

where $n_k = 10$ and $\theta_i = 2^{-i} \text{ms}^{-1}$. The autocovariance has to remain positive definite. This induces the following linear constraints:

$$\hat{c}_i = \sum_{j=1}^{n_k} \sigma_j^2 \hat{c}_i^{(j)} > 0, \quad \forall i = 1, \dots, n, \quad (16)$$

on σ_i^2 , where $\hat{c}_i^{(j)}$ are the discrete Fourier transforms of the circulant basis vectors. The optimization problem is non-concave in the subspace of σ_i^2 and multiple local maxima and saddle points can occur. Therefore, multiple initializations have to be made in order to find a potential global optimum. In general, the least-squares fit of $k(t)$ to the empirical autocovariance function Eq (14) yields a good starting point for the optimization.

The spike-rate adaptation kernel is chosen to be a linear combination of ten different alpha shapes

$$\eta(t) = \begin{cases} \sum_{i=1}^{n_\eta} w_i [\exp(-v_i t) - \exp(-\omega_i t)], & t > 0, \\ 0, & t \leq 0, \end{cases} \quad (17)$$

where we chose $n_\eta = 10$, $v_i = 2\omega_i$ and $v_i = 2^{-i} \text{ms}^{-1}$.

Since the median filter time constant is short, the voltage change around the spike can be fast, requiring flexible spike-related kernel basis. Most of this flexibility is required around $t = \delta$. Because δ is adapted, we choose a discrete parametrization which has equal flexibility from $t = 0$ up to a maximum t . In our case, this maximum is at $t = 60$ ms, and therefore our

parametrization of the spike-related kernel reads

$$\alpha(t) = \begin{cases} a_i, & \text{if } t \in [i\Delta t, (i+1)\Delta t) \\ 0, & \text{else} \end{cases} \quad (18)$$

where $a_i \in \mathbb{R}$, $i = 1, \dots, 60$ are the free parameters. Since the spike-related kernel fitting is concave, the large number of parameters does not lead to a dramatic increase of computational time. It also does not lead to overfitting, as is evidenced by the smoothness of the fitted kernels (see Figs 4 and 5 in the main text and S2–S4 Figs) and by the model comparison results (see Fig 6).

Model validation

We performed a factorial model comparison (see Fig 6) where the four factors were the presence/absence of each of the following: multiple OU components in the GP autocorrelation function (see Eq (15), as opposed to only one OU kernel with variable time-constant), the spike-related kernel α , coupling between u and s (through β) and adaptation η , which gives a total of 16 different models. We use the nomenclature that \mathcal{M}_0 is the simplest model, e.g. $\alpha = \beta = \eta = 0$ and only one OU component, having only four parameters (u_r, θ, σ and r_0). A subscript G (for GP) indicates that we use the multiple OU basis and any other subscript indicates that the corresponding parameter is adjustable in addition to the parameters already present in \mathcal{M}_0 and the parameters that are associated with the subscribed ones. E.g. $\mathcal{M}_{G\alpha}$ indicates that we use the multiple OU basis and allow a non-zero spike-related kernel and that there are now 73 parameters (δ, u_r, θ, a_i for $i = 1, \dots, 60$, and $\log r_0$). The parameter δ is optimized only for the 12 out of 16 models which depend on this parameter, i.e. that have at least $\beta \neq 0$ or $\alpha \neq 0$.

For each of the models $\mathcal{M} \in \{\mathcal{M}_0, \dots, \mathcal{M}_{G\alpha\beta\eta}\}$, we performed eight-fold cross-validation [47] on dataset \mathcal{D}_2 in order to assess the models' generalization performance. The entire dataset was cut into eight equally-sized chunks d_j , where $j = 1, \dots, 8$, each of length 15s ($n = 15000$), and six chunks of 3s d'_j , $j = 1, \dots, 6$ set aside as a test set ($n' = 3000$). Each model was then trained on seven out of eight chunks (treating them as independent samples) giving an optimal set of parameters $\Theta_j^i = \operatorname{argmax}_{\Theta} p(\{d_k, k \neq j\} | \mathcal{M}_i, \Theta)$ and training per-bin log-likelihood $p_{ij}^{\text{train}} = \frac{1}{7n} \log p(\{d_k, k \neq j\} | \mathcal{M}_i, \Theta_j^i)$. Then the validation likelihood $p_{ij}^{\text{valid}} = \frac{1}{n} \log p(d_j | \mathcal{M}_i, \Theta_j^i)$ of the left-out chunk j was evaluated. The unseen data d'_j is used for a final benchmark of model performance, where the best set of parameters is selected for each model, i.e.

$$p_{ij}^{\text{test}} = \frac{1}{n'} \max_{k=1, \dots, 8} \log p(d'_j | \mathcal{M}_i, \Theta_k^i).$$

Supporting Information

S1 Text. Supplementary Text. Information on Discrete Fourier Transforms, circulant matrices, the circulant approximation, sampling from GPs using Fast Fourier Transforms, details of the optimization scheme, gradients and Hessians used for the optimization. (PDF)

S1 Fig. Supplementary Figure. Comparison of *in vivo* and artificial data snippets for datasets \mathcal{D}_3 and \mathcal{D}_4 , analogous to Fig 3G and 3H. The scale (shown on panel D) is the same for all four panels. Vertical lines are drawn at the spiking times. (A) A 2-second sample of *in vivo* activity from dataset \mathcal{D}_3 (Zebra Finch HVC). (B) Artificial data sampled from AGAPE with parameters learned from dataset \mathcal{D}_3 . (C) A 2-second sample of *in vivo* activity from dataset \mathcal{D}_4 (mouse visual cortex). (D) Artificial data sampled from AGAPE with parameters learned from dataset \mathcal{D}_4 . (PDF)

S2 Fig. Supplementary Figure. The fitting result as a function of the parameter δ for dataset \mathcal{D}_1 , see color code next to the plot of the marginal distribution of u in the second row of the left column. The top left panel shows that the log likelihood peaks at $\delta = 18$ ms, and the bottom right panel shows the decrease of the effective coupling strength as δ increases. (EPS)

S3 Fig. Supplementary Figure. The fitting result as a function of the parameter δ for dataset \mathcal{D}_3 , see color code next to the plot of the marginal distribution of u in the second row of the left column. The top left panel shows that the log likelihood peaks at $\delta = 12$ ms, and the bottom right panel shows the decrease of the effective coupling strength as δ increases. (EPS)

S4 Fig. Supplementary Figure. The fitting result as a function of the parameter δ for dataset \mathcal{D}_4 , see color code next to the plot of the marginal distribution of u in the second row of the left column. The top left panel shows that the log likelihood peaks at $\delta = 32$ ms, and the bottom right panel shows the decrease of the effective coupling strength as δ increases. (EPS)

Acknowledgments

We would like to thank Janie Ondracek and Richard Hahnloser in Zürich, Switzerland, Daniela Vallentin and Michael Long at NYU, Bilal Haider and Matteo Carandini at UCL for kindly providing the data for this study. We also thank Máté Lengyel, Christian Pozzorini and Johanni Brea for helpful discussions.

Author Contributions

Analyzed the data: SCS. Wrote the paper: SCS JPP. Designed the software used in analysis: SCS. Developed the mathematical methods: SCS JPP. Performed the simulations: SCS.

References

1. Woody CD, Gruen E. Characterization of electrophysiological properties of intracellularly recorded neurons in the neocortex of awake cats: a comparison of the response to injected current in spike overshoot and undershoot neurons. *Brain Research*; 1978; 158(2):343–357. doi: [10.1016/0006-8993\(78\)90680-7](https://doi.org/10.1016/0006-8993(78)90680-7) PMID: [709370](https://pubmed.ncbi.nlm.nih.gov/709370/)
2. Baranyi A, Szente MB, Woody CD. Electrophysiological characterization of different types of neurons recorded in vivo in the motor cortex of the cat. ii. membrane parameters, action potentials, current-induced voltage responses and electrotonic structures. *Journal of Neurophysiology*; 1993; 69(6):1865–1879. PMID: [8350127](https://pubmed.ncbi.nlm.nih.gov/8350127/)
3. Steriade M, Timofeev I, Grenier F. Natural waking and sleep states: a view from inside neocortical neurons. *Journal of Neurophysiology*; 2001; 85(5):1969–1985. PMID: [11353014](https://pubmed.ncbi.nlm.nih.gov/11353014/)
4. Matsumura M, Cope T, Fetzi EE. Sustained excitatory synaptic input to motor cortex neurons in awake animals revealed by intracellular recording of membrane potentials. *Experimental Brain Research*; 1988; 70(3):463–469. doi: [10.1007/BF00247594](https://doi.org/10.1007/BF00247594) PMID: [3384047](https://pubmed.ncbi.nlm.nih.gov/3384047/)
5. Poulet JFA, Petersen CCH. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*; 2008; 454(7206):881–885. doi: [10.1038/nature07150](https://doi.org/10.1038/nature07150) PMID: [18633351](https://pubmed.ncbi.nlm.nih.gov/18633351/)
6. Lee AK, Manns ID, Sakmann B, Brecht M. Whole-cell recordings in freely moving rats. *Neuron*; 2006; 51(4):399–407. doi: [10.1016/j.neuron.2006.07.004](https://doi.org/10.1016/j.neuron.2006.07.004) PMID: [16908406](https://pubmed.ncbi.nlm.nih.gov/16908406/)
7. El Boustani S, Marre O, Béhuret S, Baudot P, Yger P, Bal T, et al. Network-state modulation of power-law frequency-scaling in visual cortical neurons. *PLoS Computational Biology*; 2009; 5(9):e1000519. doi: [10.1371/journal.pcbi.1000519](https://doi.org/10.1371/journal.pcbi.1000519) PMID: [19779556](https://pubmed.ncbi.nlm.nih.gov/19779556/)
8. Berkes P, Orbán G, Lengyel M, Fiser J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*; 2011; 331(6013):83–87. doi: [10.1126/science.1195870](https://doi.org/10.1126/science.1195870) PMID: [21212356](https://pubmed.ncbi.nlm.nih.gov/21212356/)

9. Lapicque L. Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *Journal de Physiologie et de Pathologie Générale*; 1907; 9(1):620–635.
10. Stein RB. The information capacity of nerve cells using a frequency code. *Biophysical journal*; 1967; 7(6):797–826. doi: [10.1016/S0006-3495\(67\)86623-2](https://doi.org/10.1016/S0006-3495(67)86623-2) PMID: [19210999](https://pubmed.ncbi.nlm.nih.gov/19210999/)
11. Latham PE, Richmond BJ, Nelson PG, Nirenberg S. Intrinsic dynamics in neuronal networks. i. theory. *Journal of Neurophysiology*; 2000; 83(2):808–827. PMID: [10669496](https://pubmed.ncbi.nlm.nih.gov/10669496/)
12. Fourcaud-Trocmé N, Hansel D, van Vreeswijk C, Brunel N. How spike generation mechanisms determine the neuronal response to fluctuating inputs. *The Journal of Neuroscience*; 2003; 23(37):11628–11640. PMID: [14684865](https://pubmed.ncbi.nlm.nih.gov/14684865/)
13. Brette R, Gerstner W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*; 2005; 94(5):3637–3642. doi: [10.1152/jn.00686.2005](https://doi.org/10.1152/jn.00686.2005) PMID: [16014787](https://pubmed.ncbi.nlm.nih.gov/16014787/)
14. Hille B. *Ion channels of excitable membranes*. Sinauer Associates Incorporated; 2001.
15. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*; 1952; 117(4):500. doi: [10.1113/jphysiol.1952.sp004764](https://doi.org/10.1113/jphysiol.1952.sp004764) PMID: [12991237](https://pubmed.ncbi.nlm.nih.gov/12991237/)
16. Gerstner W, Naud R. How good are neuron models? *Science*; 2009; 326(5951):379–380. doi: [10.1126/science.1181936](https://doi.org/10.1126/science.1181936) PMID: [19833951](https://pubmed.ncbi.nlm.nih.gov/19833951/)
17. Druckmann S, Banitt Y, Gidon A, Schürmann F, Markram H, Segev I. A novel multiple objective optimization framework for constraining conductance-based neuron models by experimental data. *Frontiers in Neuroscience*; 2007; 1(1):7–18. doi: [10.3389/neuro.01.1.1.001.2007](https://doi.org/10.3389/neuro.01.1.1.001.2007) PMID: [18982116](https://pubmed.ncbi.nlm.nih.gov/18982116/)
18. Pfister JP, Dayan P, Lengyel M. Synapses with short-term plasticity are optimal estimators of presynaptic membrane potentials. *Nature Neuroscience*; 2010; 13(10):1271–1275. doi: [10.1038/nn.2640](https://doi.org/10.1038/nn.2640) PMID: [20852625](https://pubmed.ncbi.nlm.nih.gov/20852625/)
19. Møller J, Syversveen AR, Waagepetersen RP. Log gaussian cox processes. *Scandinavian Journal of Statistics*; 1998; 25(3):451–482. doi: [10.1111/1467-9469.00115](https://doi.org/10.1111/1467-9469.00115)
20. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*; 2005; 93(2):1074–1089. doi: [10.1152/jn.00697.2004](https://doi.org/10.1152/jn.00697.2004) PMID: [15356183](https://pubmed.ncbi.nlm.nih.gov/15356183/)
21. Pillow J, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*; 2008; 454(7207):995–999. doi: [10.1038/nature07140](https://doi.org/10.1038/nature07140) PMID: [18650810](https://pubmed.ncbi.nlm.nih.gov/18650810/)
22. Paninski L, Ahmadian Y, Ferreira DG, Koyama S, Rahnama Rad K, Vidne M, et al. A new look at state-space models for neural data. *Journal of Computational Neuroscience*; 2009; 29(1–2):107–126. doi: [10.1007/s10827-009-0179-x](https://doi.org/10.1007/s10827-009-0179-x) PMID: [19649698](https://pubmed.ncbi.nlm.nih.gov/19649698/)
23. Rasmussen CE, Williams CKI. *Gaussian processes for machine learning*. The MIT Press; 2006.
24. Gerstner W, Kistler WM. *Spiking neuron models. Single Neurons, Populations, Plasticity*. Cambridge University Press; 2002.
25. Paninski L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*; 2004; 15(4):243–262.
26. Pozzorini C, Naud R, Mensi S, Gerstner W. Temporal whitening by power-law adaptation in neocortical neurons. *Nature Neuroscience*; 2013; 16(7):942–948. doi: [10.1038/nn.3431](https://doi.org/10.1038/nn.3431) PMID: [23749146](https://pubmed.ncbi.nlm.nih.gov/23749146/)
27. Buzsáki G. Theta oscillations in the hippocampus. *Neuron*; 2002; 33(3):325–340. doi: [10.1016/S0896-6273\(02\)00586-X](https://doi.org/10.1016/S0896-6273(02)00586-X) PMID: [11832222](https://pubmed.ncbi.nlm.nih.gov/11832222/)
28. Shinomoto S, Tsubo Y. Modeling spiking behavior of neurons with time-dependent poisson processes. *Physical Review E*; 2001; 64(4):041910. doi: [10.1103/PhysRevE.64.041910](https://doi.org/10.1103/PhysRevE.64.041910)
29. Lindner B, Schimansky-Geier L, Longtin A. Maximizing spike train coherence or incoherence in the leaky integrate-and-fire model. *Physical Review E*; 2002; 66(3):031916. doi: [10.1103/PhysRevE.66.031916](https://doi.org/10.1103/PhysRevE.66.031916)
30. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *The Journal of Neuroscience*; 1993; 13(1):334–350. PMID: [8423479](https://pubmed.ncbi.nlm.nih.gov/8423479/)
31. Katsaggelos AK, Lay KT. Maximum likelihood blur identification and image restoration using the em algorithm. *IEEE Transactions on Signal Processing*; 1991; 39(3):729–733. doi: [10.1109/78.80894](https://doi.org/10.1109/78.80894)
32. Bach FR, Jordan MI. Learning graphical models for stationary time series. *IEEE Transactions on Signal Processing*; 2004; 52(8):2189–2199. doi: [10.1109/TSP.2004.831032](https://doi.org/10.1109/TSP.2004.831032)
33. Gray RM. Toeplitz and circulant matrices: A review. *Foundations and Trends in Communications and Information Theory*; 2006; 2(3):155–239. doi: [10.1561/0100000006](https://doi.org/10.1561/0100000006)

34. Efron B, Hinkley DV. Assessing the accuracy of the maximum likelihood estimator: observed versus expected fisher information. *Biometrika*; 1978; 65(3):457–487. doi: [10.1093/biomet/65.3.457](https://doi.org/10.1093/biomet/65.3.457)
35. Chichilnisky EJ. A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*; 2001; 12(2):199–213.
36. Vidne M, Ahmadian Y, Shlens J, Pillow J, Kulkarni JE, Litke AM, et al. Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. *Journal of Computational Neuroscience*; 2012; 33(1):97–121. doi: [10.1007/s10827-011-0376-2](https://doi.org/10.1007/s10827-011-0376-2) PMID: [22203465](https://pubmed.ncbi.nlm.nih.gov/22203465/)
37. Jolivet R, Rauch A, Lüscher HR, Gerstner W. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *Journal of Computational Neuroscience*; 2006; 21(1):35–49. doi: [10.1007/s10827-006-7074-5](https://doi.org/10.1007/s10827-006-7074-5) PMID: [16633938](https://pubmed.ncbi.nlm.nih.gov/16633938/)
38. Cunningham JP, Byron MY, Shenoy KV, Sahani M. Inferring neural firing rates from spike trains using gaussian processes. In: *Advances in neural information processing systems*; 2007. p. 329–336.
39. Macke JH, Buesing L, Cunningham JP. Empirical models of spiking in neural populations. *Advances in Neural Information Processing Systems*; 2011; 24:1350–1358.
40. Mensi S, Naud R, Pozzorini C, Avermann M, Petersen CCH, Gerstner W. Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms. *Journal of Neurophysiology*; 2012; 107(6):1756–1775. doi: [10.1152/jn.00408.2011](https://doi.org/10.1152/jn.00408.2011) PMID: [22157113](https://pubmed.ncbi.nlm.nih.gov/22157113/)
41. Markram H, Toledo-Rodriguez M, Wang Y, Gupta A, Silberberg G, Wu C. Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*; 2004; 5(10):793–807. doi: [10.1038/nrn1519](https://doi.org/10.1038/nrn1519) PMID: [15378039](https://pubmed.ncbi.nlm.nih.gov/15378039/)
42. Pfister JP, Dayan P, Lengyel M. Know thy neighbour: A normative theory of synaptic depression. *Advances in Neural Information Processing Systems*; 2009; 22:1464–1472.
43. Long MA, Jin DZ, Fee MS. Support for a synaptic chain model of neuronal sequence generation. *Nature*; 2010; 468(7322):394–399. doi: [10.1038/nature09514](https://doi.org/10.1038/nature09514) PMID: [20972420](https://pubmed.ncbi.nlm.nih.gov/20972420/)
44. Hamaguchi K, Tschida KA, Yoon I, Donald BR, Mooney R. Auditory synapses to song premotor neurons are gated off during vocalization in zebra finches. *eLife*; 2014; 3:e01833. doi: [10.7554/eLife.01833](https://doi.org/10.7554/eLife.01833) PMID: [24550254](https://pubmed.ncbi.nlm.nih.gov/24550254/)
45. Vallentin D, Long MA. Motor origin of precise synaptic inputs onto forebrain neurons driving a skilled behavior. *The Journal of Neuroscience*; 2015; 35(1):299–307. doi: [10.1523/JNEUROSCI.3698-14.2015](https://doi.org/10.1523/JNEUROSCI.3698-14.2015) PMID: [25568122](https://pubmed.ncbi.nlm.nih.gov/25568122/)
46. Haider B, Häusser M, Carandini M. Inhibition dominates sensory responses in the awake cortex. *Nature*; 2013; 493(7430):97–100. doi: [10.1038/nature11665](https://doi.org/10.1038/nature11665) PMID: [23172139](https://pubmed.ncbi.nlm.nih.gov/23172139/)
47. Arlot S, Celisse A. A survey of cross-validation procedures for model selection. *Statistics surveys*; 2010; 4:40–79. doi: [10.1214/09-SS054](https://doi.org/10.1214/09-SS054)